

Concepts, Propositions,
and Schemata:
What Are the
Cognitive Units?¹

John R. Anderson

Carnegie-Mellon University

IT would seem that one of the surest advances of modern cognitive psychology over its behavioristic predecessor lies in its extensive postulation of mental structures and processes. Basically, cognitive psychology has reasserted the commonsense insight that if one's goal is to develop a theory of the connections between stimuli and responses, it helps to have a model of the structures and processes intervening between stimuli and responses. However, the radical behaviorist was guided by the motto "You cannot know what you cannot see." What goes on between stimulus and response takes place in the human black box which physiological psychology has not succeeded in effectively opening up. It seemed to the behaviorist that there was just no way of knowing what was going on in that box. However, at least some in cognitive psychology seem to have had greater faith in the powers of induction. Their faith is that by adding converging evidence upon converging evidence we should be able to uniquely identify what is going on in that box. If behavioral evidence is not enough, then there are additional constraints such as parsimony that can be called upon.

There are some modern cognitive psychologists who believe that the radical behaviorist was half correct, and I happen to be one who thinks so. If it ever really was in doubt, I think it can be shown (e.g., Anderson, 1976, 1978, 1979) that we cannot really know what is going on in a black box that we cannot open up. However, the radical behaviorist was also half wrong. Lack of unique identifiability is no reason not to postulate internal structure and represen-

1. Preparation of this manuscript and the research described herein is by grant BNS78-17463 from the National Science Foundation and Contract N00014-77-C-0242 from the Office of Naval Research. I am grateful to Lynne Reder for her comments and advice on this manuscript.

tations. I think that the success of cognitive psychology is a testimony to the importance of such structures in psychological theories. So, a recommendation I have made is to go ahead and get on with the business of science: Postulate some set of internal structures and processes that are consistent with the data and don't worry about unique identifiability.

It is not clear that anyone else is going to take my recommendation seriously, but I do. I face a certain dilemma in doing so: I have to choose some theory of internal structure. One important constraint on my choice is that the theory be consistent with the data. However, the conclusion about lack of unique identifiability implies that there will usually be multiple theories of internal structure and process that satisfy this constraint. Moreover, the problem of multiple theories of internal structure is not just an abstract metatheoretical dilemma. There are real options that I must choose between in trying to develop a theory. I suppose I could decide among the alternatives by tossing a coin, but I would feel unfulfilled as a scientist if I make my decisions this way. In fact I do make my choices in more principled ways than by tossing coins. I try to come up with constraints that serve to eliminate the ambiguity left by empirical data alone.

Among the additional constraints I have used are conventional ones such as parsimony, but I have found these to be not enough. In developing the ideas on representation in this paper, I have found two other bases to be particularly valuable. I would like to state them at the outset. One has to do with what I call "cognitive naturalness" (this is akin to an idea put forth by Pylyshyn, 1979, of "intended interpretations"). The idea is that while two explanations may be formally equivalent in accounting for behavior, one is a more natural interpretation of the behavior. Consider, purely for illustrative purposes, the equivalence between parallel and serial models (Townsend, 1974) for the Sternberg paradigm. The hallmark result of that paradigm is the set-size effect—that reaction time increases linearly with set size. That result is clearly more akin to a serial-search interpretation than a parallel-search interpretation. Therefore, were it not for the many other complications, the serial model would be preferable, not because of parsimony, but rather as a result of how we naturally think about processes. Thus, the issue is what is natural for our cognition as scientists. I think considerations of cognitive naturalness help one come to a conclusion about some important issues of representation.

A second constraint comes from efforts to come up with very

explicit process models. For me this has meant that my theories must be capable of implementation as computer simulations. Although any well-specified theory can be so implemented, some theories seem more satisfying when implemented than others. A computer simulation can be unsatisfying because of lack of parsimony, generality, or internal consistency, just as can nonsimulated theories. However, there is an additional potential problem with a simulation that does not really arise with a nonsimulated theory. This is efficiency of operation. To take an extreme instance, if I simulate a memory-retrieval process by a random search throughout all of long-term memory (e.g., Landauer, 1975), I would find that very unsatisfactory no matter how well it corresponded to the behavioral facts. It is probably upsetting as a theoretical proposal even without issues of simulation considered, but it is simply intolerable if one has to watch a simulation eating up CPU hours in grossly stupid search. I emphasize that the problem is not simply computation time, but computation time in conjunction with the nonoptimality of the computation. If there are obvious ways to be much more efficient, it seems too much to believe that the human brain would not have discovered these more efficient means through evolution. So, a second constraint is that the implemented theory not lead to gross inefficiencies.

In this paper I will be considering what a cognitive unit is and what evidence there is on whether propositions and schemata should be considered cognitive units. I will use as evidence both the conventional behavioral data and the less conventional criteria of cognitive naturalness and efficiency.

THE ISSUE OF THE COGNITIVE UNIT

It is typical to analyze cognition as a set of processes operating on mental data. It is also typical to think of the mental data as coming in *pieces* or *units* or *chunks*. These are the "packages" of data which the cognitive processes treat in an all-or-none manner. An important issue in cognitive psychology is determining what these units are, and considerable research can be seen as directed to this issue.

I would like to consider three levels of units that have been postulated—concepts, propositions, and schemata. Concepts are frequently taken as the unanalyzable building blocks or the primitives of cognitive representation. In semantic-network representa-

tions they are the nodes of the network. It is true that under some views (e.g., semantic features—Katz & Fodor, 1963) concepts might not be basic units. However, I will be assuming the semantic-network framework in which a concept is a basic unit or node. The meaning of the concept is attached to the concept node, but the concept can be accessed and processed as a unit independent of this attached meaning. Thus, in this framework the fact that concepts are units is given, and the question is whether propositions and schemata are units.

Propositions are composed from various configurations of the concepts. They are not arbitrary configurations, but have certain well-formedness constraints that derive from the logical notion of a proposition. That is, they are taken to be the smallest units of meaning that assert things about the world that might be reasonably judged true or false. Propositions have a "syntax" which can be used to determine what they assert about the world. There have been a number of proposals for propositional representations using semantic networks (e.g., Anderson, 1976; Anderson & Bower, 1973; Norman & Rumelhart, 1975; Schank, 1972; Simmons, 1972). In these proposals, propositions tend to be represented as configurations of nodes interconnecting more basic concepts.

A schema (Bobrow & Winograd, 1977; Rumelhart & Ortony, 1976) can be conceived of as a set of related propositions. Exactly what sets of propositions qualify as schemata is a little uncertain, but included are stereotyped sequences of events (typically called scripts—Schank & Abelson, 1977) or co-occurring descriptions describing various object concepts (most typically called frames—Minsky, 1975).

Much of our discussion of these units will be with reference to verbal experimental material. Hence, it is typical to correlate concepts with words, propositions with sentences, and schemata with passages. However, it should be kept in mind that concepts, propositions, and schemata are abstract mental constructs and that words, sentences, and passages are only operationalizations of ways to manipulate these abstract structures.

CRITERIA FOR A COGNITIVE UNIT

In defining what I mean by a cognitive unit, I need to assume a conception of cognition in which all processing takes the form of manipulating data structures in a *working memory* or an *active*

memory. One prominent instantiation of this conception of cognition is the production-system architecture (Anderson, 1976; Newell, 1973; Rychener & Newell, 1978), but this conception is more general and does not imply a commitment to production systems. For our purposes, the two important cognitive operations are the formation of links to interconnect a set of elements in working memory and the retrieval of a set of elements into working memory. A type of structure, say a proposition, is to be considered a cognitive unit if: (a) Invariably, when any elements of the unit are interassociated, all elements are interlinked. Since formation of interassociations is the means for acquiring new information, this is the criterion of *all-or-none learning*. (b) Invariably, when some of the elements of the unit are retrieved into working memory, all are. This is the criterion of *all-or-none retrieval*.

ARE PROPOSITIONS COGNITIVE UNITS?

I have in the past claimed that while the proposition is an important theoretical construct, it is not a cognitive unit (Anderson, 1976; Anderson & Bower, 1973). To see what is meant by this, let us consider the ACT (Anderson, 1976) theory of memory representation. Figure 1 illustrates the semantic-network representation that we would assign to a sentence like *Caesar crossed the Rubicon*. This structure encodes a simple proposition and interconnects *Caesar*, *crossed*, and *Rubicon*. The top node represents the proposition. It is connected to its subject *Caesar* by an *S* link and to a predicate node by a *P* link. The predicate node is connected to the relational concept *cross* by an *R* link and to the argument of the relation, *Rubicon*, by an *A* link. ACT is relatively unique among semantic-network theories in that there is a clear semantics associated with such network structures. The predicate node is taken to reflect the set of people who crossed the Rubicon, and the proposition node encodes the assertion that the subject, *Caesar*, is a member of that set.

The proposition was not a cognitive unit in ACT because it was possible for such structures to be only partially encoded or only partially retrieved. For instance, according to the ACT theory, a subject upon hearing the sentence *Caesar crossed the Rubicon* might only succeed in encoding in memory a structure like that in part (b) of Figure 1. That is, he would have failed to encode the argument link. Thus, in ACT, propositions did not satisfy the criterion of




FIGURE 1. (a) The ACT network encoding of *Caesar crossed the Rubicon*; (b) A partial encoding of (a) that might be retrieved into working memory.

all-or-none acquisition. Similarly, even if the subject succeeded in encoding the complete proposition in part (a), he might only succeed in retrieving a fragment like part (b) when he later tried to retrieve that proposition back into active memory. Thus, in ACT the proposition failed the criterion of all-or-none retrieval. In ACT, propositions were conceived of as formed from associations that were at least somewhat independent. Propositions did not have the unitary character postulated of a cognitive unit.

Research on Single-Proposition Sentences

The issue of whether propositions are learned in an all-or-none manner has received considerable research and has been the subject of some debate. I think this research can be profitably divided into research that has tried to focus on sentences that assert a single proposition versus sentences that assert multiple propositions. One of the early single-proposition studies (Anderson & Bower, 1973) concerned location-subject-verb-object sentences like *In the park the hippie touched the debutante* or, to propose a sentence for the 1980s, *In the disco the Arab approached the actress*. We prompted subjects for recall of the sentences with location (L), subject (S), verb

(V), or object (O). To each the subject could recall zero, one, two, or three other words. There are three ways to recall one or two words of the three. Thus, to each of the four probes there are eight possible patterns of recall. Table 1 provides a classification of the data from the experiment. In the table, failure to recall an element is denoted by a bar over the letter corresponding to the element. So $S\bar{V}O$ denotes that to a location cue, subject recalled subject and object, but not verb.

Table 1
Patterns of Recall to Location-Subject-Verb-Object Sentence

	Cued with							
	L		S		V		O	
3 recalled								
SVO	93	LVO	87	LSO	78	LSV	95	
2 recalled								
$S\bar{V}O$	17	$L\bar{V}O$	20	$L\bar{S}O$	18	$L\bar{S}O$	38	
$S\bar{V}O$	33	$L\bar{V}O$	36	$L\bar{S}O$	30	$L\bar{S}V$	30	
$S\bar{V}O$	33	$L\bar{V}O$	36	$L\bar{S}O$	21	$L\bar{S}V$	22	
Total:	83	Total:	92	Total:	69	Total:	90	
1 recalled								
$S\bar{V}O$	41	$L\bar{V}O$	43	$L\bar{S}O$	24	$L\bar{S}V$	41	
$S\bar{V}O$	23	$L\bar{V}O$	19	$L\bar{S}O$	16	$L\bar{S}V$	29	
$S\bar{V}O$	37	$L\bar{V}O$	36	$L\bar{S}O$	30	$L\bar{S}V$	31	
Total:	101	Total:	98	Total:	70	Total:	101	
0 recalled								
$S\bar{V}O$	461	$L\bar{V}O$	461	$L\bar{S}O$	521	$L\bar{S}V$	452	

Note that 12% of the probes resulted in recall of all three elements, 11.3% resulted in recall of two items, 12.5% in recall of one item, and 64.2% in recall of no items. Clearly, with 23.8% partial recall, subjects are not encoding and retrieving these sentences in an all-or-none manner all the time. Moreover, every possible pattern of partial recall is represented with some respectable frequency in Table 1.

There are various challenges that can be made and have been made to this conclusion of partial recall. One challenge involves presenting another set of summary statistics about Table 1. Rather than considering absolute frequency of partial recall, consider probability of recall of one word contingent on recall of another

word. For instance, contrast recall of V contingent on recall of O (denoted $P(V|O)$) with recall contingent on nonrecall of O (denoted $P(V|\bar{O})$): $P(V|O) = .637$, but $P(V|\bar{O}) = .073$. Similar patterns of contingency can be obtained by considering any other pair of words. The point is that recall of one term is much higher when conditionalized on recall of another. Thus, it might appear that propositional traces are much more unitary than would be expected by some notions of "chance."

R. C. Anderson (1974) has used such conditional-recall probabilities. He has also contrasted scoring subjects for verbatim recall versus a gist-scoring criterion. He counted subjects as correct in their recall if they recalled synonyms (*stone* for *rock*), superordinates (*clothing* for *sweater*), hyponyms (*sat* for *stayed*), and cohyponyms (*rifle* for *pistol*). Using a verbatim scoring, he found $P(V|O) = .634$ and $P(V|\bar{O}) = .084$. Using his gist-scoring criterion and scoring for guessing, he found $P(V|O) = .951$ and $P(V|\bar{O}) = .049$. So the all-or-none character of these conditional-recall probabilities becomes more extreme under a gist criterion. I have not always been able to find as extreme scores as Anderson, but I too find more extreme scores using a gist criterion. However, note that even with his extreme scores, Anderson still does find some evidence for partial recall (i.e., $P(V|O) < 1$ and $P(V|\bar{O}) > 0$).

In Anderson (1976) I examined this issue in some detail and found it to be a real hornet's nest of details, complications, and multiple interpretations. This data cannot serve as the sole basis for coming to a conclusion about the unitary character of the proposition. However, we can emerge from this data with two loose and safe generalizations: (1) When a single-proposition sentence is presented to a subject, there is some probability that the sentence will be partially recalled. (2) This partial recall is much less than would be expected under many conceptions which see the proposition as a combination of independent concepts.

Multiple-Proposition Sentences

A second source of data concerns memory for multiple-proposition sentences. For instance, Anderson and Bower (1973) considered the following type of sentence: *The Arab approached the actress who drank in the disco*. Figure 2 provides an approximate ACT representation for such sentences. Notice that elements within a single proposition tend to be closer together than ele-

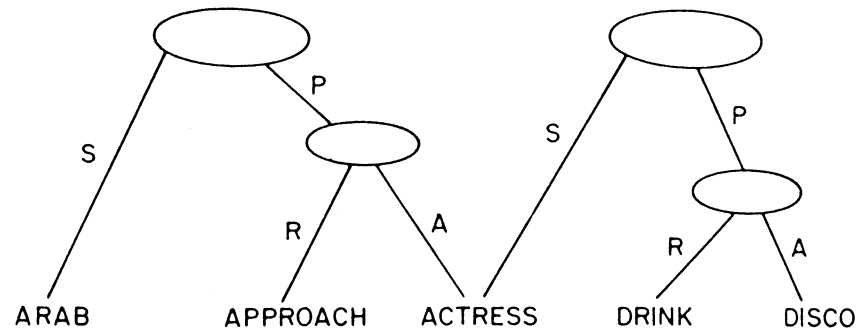


FIGURE 2. The ACT encoding for the multipositional sentence *The Arab approached the actress who drank in the disco.*

ments between propositions. Excluding *actress*, which is shared by both propositions, the mean distance between concepts within a proposition is 2.5 links, while the mean between-proposition distance is 5.5 links. So, we would expect even on the ACT account that probability of recalling an element from one proposition given an element from the same proposition would be higher than given an element from another proposition. Anderson and Bower (1973) tried to fit the HAM theory (which involved representations like Figure 2) to the recall of such sentences. We found systematic misfits of the theory in the direction of there being greater tendency to recall an element from the same proposition as the cue. That is, even though the theory predicted within-proposition contingencies, it does not predict enough. Although it is undoubtedly the case that some process assumptions could be derived that would satisfactorily predict the data given the representation in Figure 2, this systematic failure of the ACT theory can be taken as some evidence for the unitary character of the proposition. Goetz, Anderson, and Schallert (1979) and Kintsch (1974) have reported similar studies of multiple-proposition sentences and have come to similar conclusions.

Ratcliff and McKoon (1978) report a different methodology for looking at memory for multiple-proposition sentences. They had subjects commit to memory sentences like *Geese crossed the horizon as the wind shuffled the clouds*. Since subjects memorized these sentences more or less perfectly, probability of recall was not the measure of interest. Rather, the subjects saw a series of words, and they were to respond *yes* or *no* whether each word was in one of the sentences they had studied. Ratcliff and McKoon were interested in the speed of this decision.

They contrasted three conditions under which the subject might have to decide whether he had studied a word like *horizon*. In the within-proposition condition, the preceding word to be judged had come from the same proposition—thus, the preceding word might be *geese*. In the between-proposition condition, *horizon* would be preceded by a word from the same sentence but from a different proposition—thus, the preceding word might be *wind*. In the unprimed condition the word would be preceded by a word from a different sentence. The mean response times in the three conditions are 709 msec for within-proposition priming, 752 msec for between-proposition priming, and 847 msec for the unprimed condition. So, there is a clear advantage for within-proposition cueing, although there is also a considerable between-proposition advantage among the propositions connected in a single sentence.

It would be wrong to conclude that any of these studies on multiple-proposition sentences clearly establishes the unitary character of the proposition. As we noted with respect to Figure 2, representational theories like ACT or HAM assume that the elements within a proposition are more closely connected than the elements between propositions. It is true that some of the within-proposition contingencies are stronger than predicted in the HAM model, but there are other non-HAM process assumptions that could be marshaled to predict this degree of contingency.

Conclusions about the Proposition as a Cognitive Unit

The fact of fragmentary recall of propositions, even if not a high-frequency occurrence, is the major argument against concluding that the proposition is a cognitive unit. The HAM theory makes certain predictions about fragmentary recall of propositions. For instance, verb and object are closer together in the network representation and should tend to be recalled together. Some of the early research had indicated that this fragmentary recall conformed to predictions derived from the HAM structure. However, subsequent research failed to support such conclusions. This research has found some systematic effects in fragmentary recall of propositions, but these turn out to be better predicted by closeness in the surface sentence than closeness in a deeper semantic representation. That is, the fragments that tend to be recalled to-

gether are close together in the surface sentence (see Anderson & Bower, 1973, for detailed discussion).

The above research suggests that when we observe fragmentary recall, we are not observing recall of single propositions, but recall of surface strings of words. Anderson and Paulson (1977), using verification times, provided evidence that memory for a sentence can rest either on an abstract proposition or a surface string. The resolution to the issue of partial recall of propositions may lie in the distinction between propositional and word-string representations for sentences. Perhaps the propositional representations are stored and retrieved in an all-or-none manner, while the word-linked-to-word sentence representations have the fragmentary character. This proposal would serve to account for the empirical data. The empirical data is sufficiently complex now that I do not think that there is a simpler explanation than this two-trace theory.

The strongest motivation for the two-trace proposal is not based on empirical data, but rather on five years of experience with a simulation of ACT that has had to work with representations of propositions or of word strings. We have found it extremely frustrating to work with a system where basic propositions can be partially encoded or partially retrieved. A partial proposition is virtually never of any use to the system, and effort that goes into partial encoding or retrieval is simply wasted. It makes the system much less efficient than it needs to be. Essentially, what was always happening is that the system would be looking for a proposition to guide its processing, retrieve only part of the required proposition, and so be delayed in its processing or, worse yet, be misdirected. The inefficiency associated with partial propositions is not just a consequence of ACT's implementation on a computer, but rather reflects a fundamental incompatibility between partial processing of a proposition and the all-or-none significance of a proposition in information processing. It is rarely the case that a partial proposition provides any guide to information processing. (What can we do with the partial proposition *Carter defeated*, and contrast this with the implications of *Carter defeated Kennedy!*) On the other hand, fragmentary processing of word strings is a necessity given the recursive or iterative character of the parsing that underlies sentence comprehension. One has to interpret word-string fragments of a sentence and then concatenate the interpretations of these fragments to come to the sentence meaning.

In summary, I think that the conjunction of the complex empiri-

cal literature with some clearer considerations about processing efficiency supports the conclusion that propositions are treated as cognitive units. It is also the case that the criterion of cognitive naturalness, discussed earlier, points in the same direction. Given the relative dominance of all-or-none recall, it seems more natural to think of the test results as being produced by a basically all-or-none process with a set of perturbing processes producing occasional partial recall, than to think of them as being produced by a basically partial process with various factors producing an enhancement of all-or-none recall. It needs to be emphasized that this is not a matter of parsimony.

ARE SCHEMATA COGNITIVE UNITS?

Schemata can be conceived of as being composed of propositions, just as propositions can be conceived of as being composed of concepts. For instance, consider one of the 16-sentence stories that we have used in some of our research:

1. Willa received a telephone message.
2. Willa was told that her father was dying.
3. Willa's father lived in San Francisco.
4. Willa had to get to San Francisco quickly.
5. Willa called up for a taxi.
6. Willa went to Kennedy Airport.
7. Willa purchased a ticket for San Francisco.
8. Willa got on board the plane.
9. Willa flew across the country.
10. Willa ordered four drinks on the plane.
11. Willa arrived in San Francisco Airport.
12. Willa was dizzy when she arrived in San Francisco.
13. Willa did not know her way around the Bay Area.
14. Willa was incoherent when she asked for directions.
15. Willa could not find her father's hospital.
16. Willa cried in the streets of San Francisco.

The middle six sentences (6-11) correspond to what most people would consider a script or schema about an airplane trip. They reflect a fairly stereotypic sequence of events on an airplane trip. (These six sentences were adapted from some unpublished re-

search of Abelson and Reder.) The interesting question is whether these middle sentences form a "unit" in any interesting sense.

Figure 3 gives a network representation of the plausible interconnections among the 16 sentences. For instance, consider sentence (9) about Willa flying across the country. It has connections of temporal adjacency with the immediately preceding and following sentences. It is connected with sentence (7) which indicates the destination of the flight and with sentence (11) which indicates the successful completion of the flight. Although this representation does not assign any special unit status to the plane-trip episode, the network does illustrate the fact that the material tends to be more interconnected within a schema than between schemata. For instance, each schema sentence in Figure 3 has a mean of 1.0 connections to nonschema sentences, but 3.3 connections to schema sentences.




FIGURE 3. A diagrammatic representation of the relation among the sentences in the story about Willa.

This story representation has an important feature in common with our earlier ACT network representation of sentences (refer back to Figure 2). That is, in the story representation the elements composing a particular schema are more closely interassociated, just as in a sentence the concepts composing a proposition are more closely interassociated. Thus, even if we did not ascribe any special status to the schema, we would expect to see schema-defined contingency in story memory. Analogous to the propositional case, we can ask whether the contingency in schema recall can be explained solely in terms of this greater density of interassociations, or whether we will have to propose that schemata form special units. That is, are schemata encoded and retrieved in separate propositional units, or are they encoded and retrieved in an all-or-none manner?

General Schemata versus Instantiated Schemata

There is a fundamental ambiguity that lies in the use of schemata to refer to memory objects. On one hand, schemata can refer to general stereotyped patterns of propositions—for instance a general statement of what a dog looks like or what happens in a restaurant. This is probably their primary use. In such a case they are part of one's general semantic knowledge and are not the sorts of things that are generally learned in a memory experiment. Thus, the question of all-or-none learning does not arise experimentally. It is natural, however, to suppose that they gradually evolve through experience. The question of all-or-none retrieval is relevant. If a particular instantiation of such a schema is presented as to-be-learned material, and if the schema is retrieved to be put in correspondence with this experimental material, are the elements of that schema retrieved as a unit or individually?

The second use of schema is to refer to what is deposited in memory as a function of making the correspondence between presented material and a schema. According to schema theory this experimental material has somehow been organized to reflect the structure of the schema. For this use of the term *schema*, we can inquire whether it meets both defining features of a cognitive unit. Is the material within a schema interassociated in an all-or-none manner? Are its elements retrieved in an all-or-none manner? In this paper I am mainly concerned with this second use of schema and whether it can be considered a cognitive unit. How-

ever, I will also suggest that there might not be fundamental differences between these two types of schemata. Therefore, conclusions about the status of one sense of schema might generalize to conclusions about the other sense.

It would seem that typical discussions of schemata imply at least a partial commitment to their special status as cognitive units. This seems clearest in the writing of Schank and Abelson on scripts, which is the most explicit instantiation of a schema system. Mention of certain "header terms" like *restaurant* for the restaurant script caused the script to be invoked. When invoked, all the information in the script is available for processing a text. This would appear to be like having the script recalled in an all-or-none manner into working memory. While schemata appear to have a property of all-or-none retrieval, it does not seem that schema theorists are committed to all-or-none encoding of elements that conform to a schema. Schema theorists have not been explicit about what processes govern probability of encoding a fact except to suggest that facts which do not fit a script are less likely (or more likely—Schank, 1975) to be encoded. It is perhaps most in keeping with existing discussion of schemata to propose that facts which fit in a schema have an increased probability of recall, rather than to propose that they are remembered in an all-or-none manner. These two proposals can amount to quite different things, as a little thought should confirm.

Research on Schemata

There has been research showing distortions in memory such that subjects falsely think they studied things normatively part of a script and such that they transform other facts to fit the normative prescriptions of the script (Bower, Black, & Turner, 1979; Mandler & Johnson, 1977; Rumelhart, 1977; Thorndyke, 1977). While such studies may show the importance of schemata as explanatory concepts, they are not directly relevant to assessing whether schemata are cognitive units.

Black and Bower (1979) and Owens, Bower, and Black (1979) looked at the recall of multiepisode stories where each episode might be conceived of as a schema. They found that subjects recall facts from only some of the episodes, and from each of these they can retrieve a fair number of facts. This result is reminiscent of the some-or-none recall of instances from a category in a free-recall list (Cohen, 1963, 1966).

It is worth reviewing the basic phenomenon of some-or-none recall of items from a categorized list. If one plots the distribution of number of instances recalled from n -member categories, that distribution conforms to a binomial distribution except that there are too many cases of zero items recalled. It is as if the subject has one probability a of accessing the category and another probability b of retrieving instances for an accessed category. The $1-a$ probability of failing to access the category produces the concentration of zero recalls from a category, and the b probability, once in the category, produces the otherwise good fit to a binomial distribution.

Analogous to the situation of a categorized list, it might appear in script recall that there is one probability of retrieving the schema and then a second probability of retrieving the individual propositions organized by a schema. This implies, which is empirically the case, that all the facts are not recalled from a schema. Rather recall is some-or-none. However, this leaves open the issue of whether the failure of the "all" part of all-or-none recall occurs at encoding or at retrieval or at both. It should be kept in mind that there does not seem to be a commitment in current expositions of schema theory to all-or-none encoding. It might be that the information organized by a schema is partially encoded at study and that this partial encoding is retrieved in an all-or-none manner. This would produce the observed phenomenon of some-or-none recall.

An Experiment Looking for Schema-Units

We have recently performed a further study to determine whether there is all-or-none retrieval of schema-organized information. This study used a priming manipulation in a reaction-time paradigm. The study is analogous to the one by Ratcliff and McKoon described earlier. Such a methodology seems more appropriate than simple-recall studies for getting at the retrieval dynamics implied in the issue of all-or-none retrieval.

Our study used 16-sentence passages such as the Willa passage, where the middle 6 sentences corresponded to well-defined scripts (as determined in unpublished research of Abelson and Reder). We reasoned that if these scripts were to have a unitary quality, there should be a special facility for the sentences in one script to prime other sentences in that same script. We had subjects commit

to memory six such stories. Then we used a sentence-recognition paradigm in which subjects had to decide whether a particular sentence came from one of the stories they had studied. Foils were created by re-pairing subjects with predicates. On most of the trials we avoided repetitions of sentences from the same story, but we created a few (10% of the trials) critical repetitions. These could involve pairs of sentences from inside the script or one sentence from outside the script and one from within. Always the second, primed sentences were from within the script. In fact these critical second sentences were the same in the outside-priming and inside-priming conditions. We were interested in whether there would be special priming and hence a special recognition advantage if the first sentence came from inside the same script as the second sentence. A special advantage for inside sentences seems implied from the viewpoint that the script is a cognitive unit. The principle of all-or-none retrieval implies that bringing one script sentence in working memory should bring the other script sentences in.

However, a priming advantage for within pairs might occur even if there were not special unitizing of the script. This can be seen from inspecting the greater interconnectedness of within-script sentences in Figure 3. The sentences within the script tend to be more related to one another, and priming might just reflect this greater relatedness. Therefore we had subjects rate pairs of sentences, either both from within the script or one from within and one from without. We chose pairs of sentences for our inside- and outside-priming conditions that were approximately equal in relatedness. Moreover, to assess whether there was an effect of relatedness, we divided the inside and outside pairs into two groups—those that received high relatedness ratings and those that received medium relatedness ratings.

Table 2 presents the recognition times for the second sentence in primed pairs, classified according to whether they were instances of outside or inside priming and according to the rated degree of relatedness for the pair. It is clear from Table 2 that there is little effect of either variable—certainly no significant effect. These times should be compared with subjects' response times to these same sentences when they were preceded by a sentence from another story—1,240 msec with a 6.9% error rate. Thus there is clearly priming relative to this control, but there is no effect of either priming variable in Table 2.

Our stories were deliberately constructed in this first experi-

Table 2

Recognition Times in Msec (and Error Rates in Parentheses) to Primed Members of Sentence Pairs

		Type		
		Inside	Outside	Average
Degree of Relatedness	Hi	1078 (.028)	1109 (.038)	1094 (.033)
	Med	1071 (.027)	1087 (.035)	1079 (.031)
Average		1075 (.028)	1098 (.037)	1086 (.032)

ment to be thematically integrated. That is, the six-sentence script episode fit in as an integral part of the full story. This was to create examples of between-script pairs that were as related as within-script pairs. It seems like a fair test for whether there were unique properties associated with script processing. The story as a whole was not a stereotypic sequence of events, whereas the subset of script sentences was. If stereotypic sequences are specially processed as units, there should be special priming of the sentences within a script over and above sentences which are just thematically related.

The conclusion of this research seems to be that there is a special facilitating effect of thematic relatedness, but no special effect of the kind of stereotypic relatedness we associate with scripts. Note, however, there was not an effect of degree of relatedness. One might conclude that the thematic relatedness of the story had made the whole story a cognitive unit. To test this interpretation, we ran a second experiment in which the whole story was not thematically integrated.

An Experiment on Disjointed Stories

An example of the kind of stories used in this research follows:

1. Willa went to the sink to brush her teeth.
2. Willa picked up the tube of toothpaste.
3. Willa put some pepsodent on her toothbrush.
4. Willa cleaned her teeth in front and back.
5. Willa rinsed her toothbrush in cold water.

6. Willa went to Kennedy Airport.
7. Willa purchased a ticket for San Francisco.
8. Willa got on board the plane.
9. Willa flew across the country.
10. Willa ordered four drinks on the plane.
11. Willa arrived in San Francisco airport.
12. Willa entered the shoe store.
13. Willa found a shoe salesman.
14. Willa tried on a pair of shoes.
15. Willa said she would buy the shoes.
16. Willa took the boxed shoes home.

Note that the first five sentences are integrated, the next six are, and the last five are; but there are major "cognitive jerks" in the transition from one set to another. The central six sentences were the same ones as used in the previous experiment. We also used the same within-between manipulation in this experiment as previously. That is, in the critical pairs either both sentences came from within the critical six (within-primary condition), or the first came from without and the second from within (between-primary condition). As before, physical distance between the pairs, in terms of number of intervening sentences, was the same in both conditions.

We expected that if thematic relatedness was sufficient to create a cognitive unit, we should find priming within the six-sentence script and not between. The results were 1,134 msec (1.2% errors) for within-script priming, 1,240 msec (5.3% errors) for between-script priming, and 1,390 msec (6.4% errors) for the control case when the target sentence was preceded with a test item taken from a different story. Clearly in this experiment we have evidence for within-script priming. The between-script condition is still faster than the control condition. It is a little hard to know why the between-script condition remains better than the control. It may just reflect the benefit of repeating the person's name across the pair of sentences. In any case, we have obtained clear evidence for within-script priming when that script is thematically distinct from the rest of the passage. This reinforces the conclusion from the first experiment that what is important in determining the cognitive unit is thematic relatedness and not scriptal character per se.

Thus, it seems that we have evidence for cognitive units larger than the proposition. If we are to call these cognitive units schemata, we will have to change somewhat our interpretation of

schemata away from the notion of a stereotypic set of facts or events to a notion of a thematically related set. This also carries with it the requirement that we become more precise about what we mean by "thematically related." I will attempt to do this after discussing some further research.

Cognitive Units and Interference

Another way to get at the issue of whether schemata are cognitive units involves the study of retrieval interference. The phenomenon of retrieval interference refers to the fact that as more items are associated to the same element, there is a loss in the facility to retrieve any of these items. This has been classically shown in a paired-associate paradigm where learning multiple responses to the same stimulus causes the responses to interfere with each other.

More recently the interference phenomenon has been studied in a sentence-recognition paradigm (e.g., Anderson, 1974, 1976; Hayes-Roth, 1977; Thorndyke & Bower, 1974). Subjects are asked to commit to memory a set of facts like: *The lawyer bought a pair of Adidas. The lawyer entered the restaurant. The doctor waited for the train.* Subjects are then asked to judge whether they have studied such sentences when they are mixed up with foils like *The lawyer waited for the train.* The more facts that a subject learns about a particular concept like *lawyer*, the slower he is to recognize any fact he has studied or to reject a fact he has not studied. This result is referred to as the *fan effect*. In the network representation of this information, more facts learned about a concept like *lawyer* means a greater fan of links leading out from the concept. Assuming that the rate at which activation spreads from the concept down the links is an inverse function of this fan, and assuming that retrieval depends on activating this structure, then recognition time will depend on fan.

Recently Smith, Adams, and Schorr (1978) uncovered an interesting situation where this fan effect does not hold. They presented subjects with a set of facts which could be unified into a schema. An example would be: *The lawyer christened the ship. The lawyer broke the bottle. The lawyer did not delay the trip.* As long as the facts could be organized into such a schema, there was no effect of number of facts on recognition time. Smith et al. suggest that such sets of facts are organized into script units. They suggest that while a script unit may take a long time to access, once it is

accessed, its individual members can be scanned very rapidly. Thus, they are basically taking a variant of the schema-as-cognitive-unit approach. Translated into the terminology of activation or working memory, their assumption amounts to the claim that the time to activate a script or to read it into working memory is little affected by the number of facts contained in that script.

Some of our recent research (Reder & Anderson, 1980) has complicated in interesting ways the account of the Smith et al. result. It turns out that the Smith et al. result depends on the nature of the foils used. In various experiments Smith et al. created their foils by combining sentence subjects like *lawyer* with predicates that had not been studied of *lawyer*. They constructed their foils in two ways: (1) The predicate had been studied with some other character, but was not related to the script studied of the tested person. Thus, the subject might learn ship-christening facts about the lawyer and, say, restaurant facts about the dentist such as *The dentist ate the hamburger*. A foil would be *The lawyer ate the hamburger*. (2) The predicates had not been studied about another person, but were related to the script studied about the person. These predicates were derived by adding an extra word to a predicate that had been studied. So, whereas the subject had studied *The lawyer broke the bottle*, he would be tested with *The lawyer broke the champagne bottle*.

In our experiment we also used foils where (3) the predicate had both been studied about another person and was related to the original script. So the subject might have studied ship-christening facts about the doctor as well as the lawyer. For the doctor one of the facts might have been *The doctor waved good-bye*. The foil for lawyer would be *The lawyer waved good-bye*. We found that when we used foils like (3) the fan effect reemerged, whereas we were able to replicate Smith's result of a reduced fan effect with foils like (1). Subjects are also much faster at rejecting foils like (1) than foils like (3). It seems that subjects were able to use the lack of a relation to the script in foils like (1) to reject the foils quickly. Similarly, we suspect that they used the simple presence of a script relation to accept targets, thus avoiding the need for a search of the script. We ~~did not try foils like (2), but we suspect that subjects used the~~ lexical unfamiliarity of terms like *champagne* to reject these. Subjects appear to engage in a search of the script only with foils like (3) which are related to the studied script and which have no unfamiliar terms.

To summarize this aspect of our research, it seems that a revised

interpretation of the Smith et al. research is necessary. It appears that the individual propositions can be unified into a schema, and that subjects can treat this schema in an unanalyzed fashion, provided they do not have to make fine discriminations between what they actually studied versus what is only related to the schema. However, if we force subjects to make such discriminations, then they must analyze the contents of their schema and the fan effect re-emerges.

There is another aspect of our experiment which is very relevant to the issue of schemata as cognitive units. We manipulated whether subjects studied one or two schematically organized clusters of facts, and the number of facts in each cluster. Thus, in the two-schema condition, subjects might study facts about the lawyer in the restaurant and the lawyer taking a train. We were interested in what the effects might be of learning a second schema, and of the number of facts in that schema, on time to verify facts from the first schema. Our results were quite clear: (1) Subjects were slowed by the presence of a second schema. (2) Subjects were not affected by the number of facts in that second schema. These effects held whether the foils were unrelated (type 1) or related (type 3). Both major results of our experiment are interesting. The first implies that while the facts in a schema do not interfere with each other at least under some circumstances, they can be interfered with by the presence of other unrelated facts. The second result implies that the interfering effect of these other unrelated facts can be minimized if they can be encapsulated into a schema themselves.

Figure 4 illustrates a slight adaptation of the knowledge structure proposed in Reder and Anderson. There is a node in memory corresponding to the lawyer. It is connected to what we called "subnodes" to represent the train and restaurant events. Directly attached to each subnode is an indicator of the schema it represents. Also attached to each subnode are the proposition nodes corresponding to the individual facts organized by the schema. Each of these proposition nodes is then connected to the individual concepts organized by the proposition. (Each of these propositions should include reference to the lawyer, but this is omitted for sake of simplicity.) What we have here, then, is a hierarchy of the sort seen many times in cognitive psychology. The lawyer can be seen as a cognitive unit that serves to organize its individual schemata. Each schema is a cognitive unit that or-

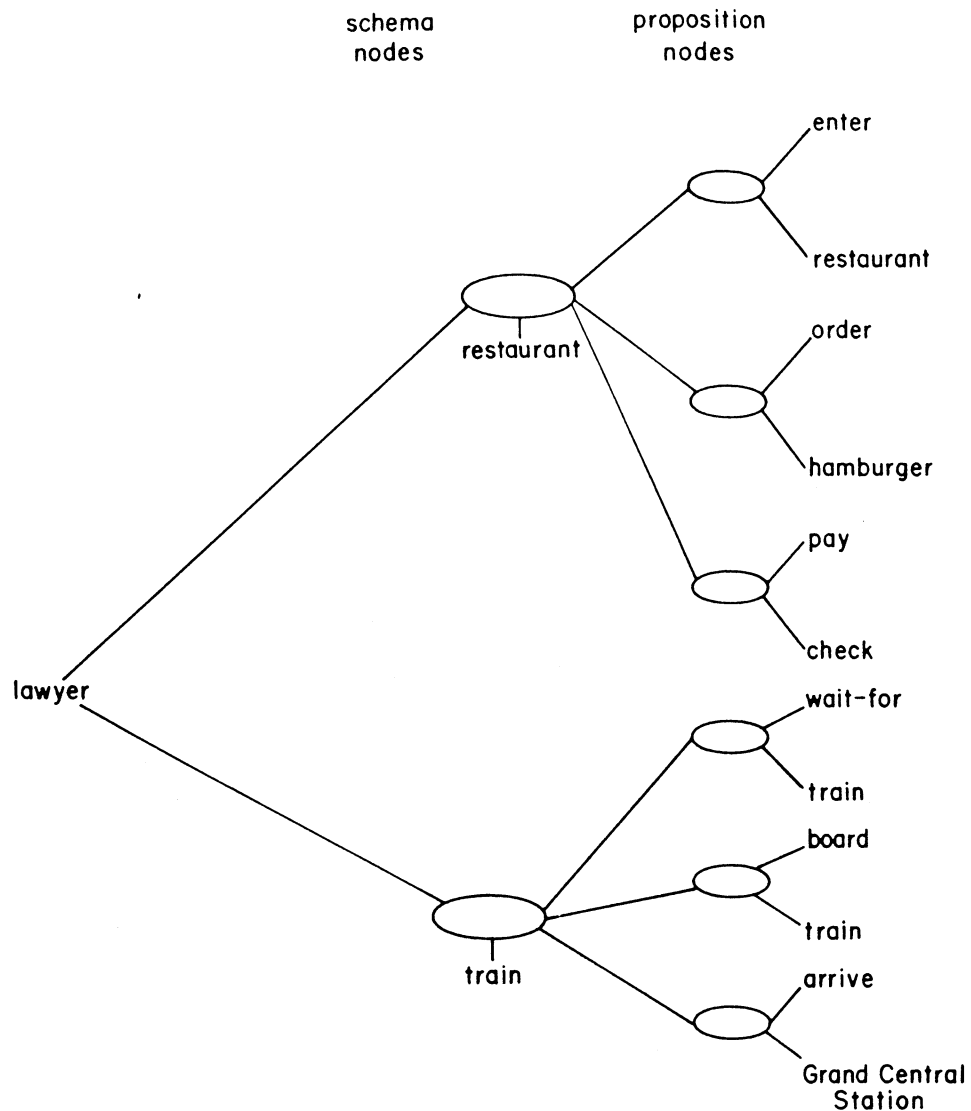


FIGURE 4. A representation of the hierarchical structure of schemata, propositions, and concepts attached to an individual in the Reder and Anderson experiment.

ganizes individual propositions. Each proposition is a cognitive unit that organizes individual concepts.

Faced with unrelated foils, subjects needed only to retrieve the schema nodes with their tags. They did not have to unpack each schema into its propositions. It is hypothesized that subjects can accept or reject probes by simply deciding if the probes were from the appropriate schema. This means that with unrelated foils, de-

cision time should be a function of number of schemata but not number of propositions within a schema. This is just what we found. With related foils, subjects would have to inspect the individual propositions in the relevant schema. This would mean that decision time should vary with number of schemata, number of facts in the relevant schema, but not number of facts in irrelevant schemata. Again, this is just what we found.

An important issue in this proposal is how subjects decide whether a schema node is appropriate or not without inspecting the contents of the schema. There are two solutions that we have proposed. The first is that subjects retrieve a schema label from the predicate of a probe and match this to the labels stored with the schema nodes attached to the individual. There would be some implementation difficulties in achieving this, and even then the schema would have limited applicability. A basic problem is that a predicate may participate in many schemata, and it is not clear how to select the correct one without first considering all the information attached to the individual. A more reasonable solution in our view simply uses level of activation. The idea is that activation spreading from the individual node (e.g., lawyer in Figure 4) will quickly reach the schema node. Similarly, the predicates are strongly associated to the schema labels by prior knowledge. This is not shown in Figure 4, but there are strong preexperimental associations between propositions or predicates like *pay check* and schema labels like *restaurant*. Thus, activation will quickly converge on the schema node from the predicate. Because of this intersection of activation between the individual and the predicate, the node for the correct schema will be more active, allowing it to be selected.

Other Considerations about Schemata

So far I have presented some memory data in favor of the conclusion that schemata are cognitive units. For instance, we have shown in the experiment on disjointed stories that we can get within-script priming when the script has a sharp thematic contrast with surrounding information. The fact that the advantage of within-script priming (or the deficit of between-script priming) disappears when the thematic boundary is taken away, points to the fact that we need to liberalize the notion of a schema considerably over that of a script. With this liberalization the results are

consistent with the schema-as-cognitive-unit view. Again the results of Smith et al. and of Reder and Anderson are nicely explained within the view that schemata are cognitive units.

While these results are conveniently explained in the schema-as-unit approach, they do not argue uniquely for this explanation. Indeed, in Reder and Anderson our analysis of these experiments did not invoke a units concept. My belief that schemata are cognitive units, as is the case for propositions, is bolstered by considerations of efficiency and cognitive naturalness.

Just as our efforts at simulation led us to dissatisfaction with the fragmented encoding of propositions, so too our efforts at simulation have led us to be dissatisfied with our inability to treat sets of propositions as units. Sometimes it proves desirable, and on other occasions it is essential to treat sets of propositions as units distinct from other related propositions. These are occasions when we need to treat sets of propositions as *patterns*. For instance, the set of propositions about a person's face may need to be treated as a pattern distinct from other propositions about the person. To take a more abstract example, the set of facts given in a geometry-proof problem needs to be treated as a pattern distinct from other related information in searching memory for solutions to similar proofs.

In the ACT system of 1976 there was already the facility to treat as a unit the propositions that made up the condition pattern of a production. These productions, however, were part of ACT's procedural knowledge of how to do things. What we are now finding is the need to deal with patterns in ACT's declarative-memory component.

The essence of a pattern is that it is to be put in correspondence with another object. By declaring a set of propositions a unit, we partition them as the elements for which a correspondence is to be found. This is the convenience of unitization. When patterns involve variables, partitioning is no longer a matter of computational convenience but rather is a necessity. The clearest argument for this point was developed by Hendrix (1975) in his discussion of partitioned semantic networks. For instance, consider the following information. *Labradors are large black dogs with floppy ears.* Figure 5 illustrates the network connections among the basic propositions that encode the sentence above. Here we have a set of propositions about a variable *X* which serves to stand for the defined labrador. The structure in Figure 5 might be translated, "If *X* is a labrador, then *X* is large, *X* is a dog, *X* is black, and *X* has *Y*, and *Y* are floppy

matched. Thus, it is hard to argue that there would be any disaster if only part of a pattern were retrieved from memory. All one can say is that partial encoding or retrieval would further degrade the reliability of the matching process.

It is at this point where cognitive naturalness comes in as a criterion for judging a representation. It is possible, as suggested above, that structures which need to be logically treated as units in pattern-matching will still be retrieved in a piecemeal manner. However, it is clearly unnatural to have this dichotomy between the pattern-matching processes and the retrieval processes. Again, while it is possible to come up with explanations that account for the memory phenomena that were presented as evidence for units interpretation without invoking the concept of a cognitive unit, there is something unnatural about explaining data suggestive of units with piecemeal mechanisms.

In conclusion, I think that the issue of whether schemata should be considered cognitive units is currently less decidable than the propositional case. We need more data and more quasi-logical arguments before enough evidence will be available on which to make a clear decision. However, I also think that the current state of evidence favors the cognitive-unit interpretation.

COGNITIVE UNITS: THEIR ENCODING AND RETRIEVAL

So, I have come to tentative answers to the questions set out at the beginning of the paper. It seems that propositions and schemata can be cognitive units by the encoding and retrieval criteria set forth at the beginning of the paper. I think that with more explicitly stated assumptions about encoding and retrieval, we can account for the apparent empirical contradictions to this conclusion. However, before developing the more explicit assumptions, I would also like to suggest that propositions and schemata not be seen as two different species of cognitive units, but rather that they be viewed as two manifestations of the same species. I do not think there will prove to be any useful distinction between propositions and schemata beyond size. The basic idea is that when elements in working memory can be put into correspondence with the elements of an existing knowledge structure, the working-memory elements will be unitized. The concepts in a proposition form a unit because

they can be put into correspondence with various well-worn cognitive schemata such as *agent performing action on object* or *object in location*. Similarly, sets of propositions can be unitized when they can be put into correspondence with a cognitive schemata. Thus, I am treating cognitive units to have the same range of denotation as is exemplified by schemata in a system like KRL (Bobrow & Winograd, 1977). I will return later to a fuller discussion of what is meant by putting the elements in working memory in correspondence with the elements of an existing cognitive unit.

The idea that there is a general cognitive-unit structure that spans concepts, propositions, schemata, and other structures as special cases is appealing, of course, because of its generality. The generality of the idea is part of what leads me to venture a hypothesis about schemata in the face of insufficient evidence.

Encoding Assumptions

Now I would like to spell out more carefully what is meant by unitization and what the encoding and retrieval consequences are of unitization. Unitization occurs when a set of elements in working memory can be put into correspondence with an existing knowledge structure. Unitization involves creating a new node, let's call it the *chunk* node, and adding links to interconnect the previously unconnected working-memory elements. Also, an *index* tag is added to this chunk. This index points to the prior-knowledge structure with which it had been put into correspondence. The index provides a way to access the content of a chunk without actually inspecting the chunk. We discussed two ways we could do this with respect to Figure 4 for the Reder and Anderson experiment.

It is useful to contrast the encodings of elements that form cognitive units with the encodings of elements that do not. If elements cannot be put into correspondence with a cognitive unit, they still can be linked together to a node. However, each of the links must be formed independently. This gives us the contrast between all-or-none encoding and partial encoding. Secondly, there is no index that can be attached to guide access to these elements. This means, for instance, that one cannot reject these elements as a unit, and so one cannot reduce interference as in the Reder and Anderson study.

It should be emphasized that the contrast between the unit and nonunit is one of rate of encoding as well as all-or-none encoding.

The amount of effort (number of separate acts of association or time to interassociate) required in the case of a nonunit is linear with the number of elements being associated. In the case of the unit the effort is independent of the number of elements being associated. This accounts for the heightened level of recall typically associated with elements that organize themselves into well-formed cognitive units. It is interesting to speculate about the kind of mechanism that would find it easier to create structure given a corresponding structure. If we think of the links to be formed as complex and requiring some effort to create, then something of an explanation can be offered out of analogy to computer systems. As a rule it is much easier to create a data structure in a computer if the structure can be simply copied from an existing structure rather than if it has to be created anew. Perhaps the function of a correspondence is to enable the human memory system to create the new structure by the analog of a copy operation rather than having to perform the analog of creating structure by original computation.

One empirical difficulty with the all-or-none encoding assumption was the apparent result that multiproposition units in texts or stories are only partially recalled. I think this apparent contradiction can be explained away by assuming a difference between the operational definition of units and what is a cognitive unit for the subject. That is, there may be all-or-none encoding of multiproposition schemata, but these schemata do not always correspond to what the experimenter designates as units. A definition of what is and what is not a schema has not been as sharp as the definition of a proposition. Moreover, if we expect a subject's schemata to depend on his experience—what he has learned to be a unit—then we might well expect subjects to differ among themselves as to what the boundaries are for their schema units.

In particular I think there is very good reason to believe that the true cognitive schemata are much smaller than many of the schemata suggested in recent theoretical proposals. The size of any structure that can be unitized is limited by the number of propositions that can be kept active in working memory. It is the conventional wisdom to assign a relatively small number to working-memory capacity (e.g., four elements—Broadbent, 1975). This is much smaller than the number of elements that are listed in many schemata. The only way the large schemata can be organized is to have hierarchies of schemata with higher level schemata organizing lower level schemata. Indeed, this is what Schank and Abelson seem to be doing with their scenes within scripts. The hierarchical

structure of many text units is explicitly recognized in the story grammars. It has also been very clearly acknowledged in some of the other discussions of schema theory (e.g., Bobrow & Winograd, 1977; Rumelhart & Ortony, 1976). Of course, if a text must be encoded into a hierarchy of cognitive units, it is possible to have partial encoding of the whole text, even if each of the cognitive units in the text is encoded in an all-or-none manner.

Retrieval Assumptions

When the cognitive unit is activated, the chunk node, its index, and its elements become available in working memory. This could be referred to as the *expansion* of the cognitive unit. The elements of a cognitive unit may themselves be cognitive units. However, when a cognitive unit is activated and expanded in working memory, the elements of the cognitive unit are not also expanded. Activation must spread from the chunk node to the element in order that the element be expanded. This means that, in the Smith et al. and the Reder and Anderson experiments, subjects were able to quickly identify the theme of their cognitive units but not their specific contents. The expansion of a cognitive unit is independent of the number of elements in that unit. The effect of number of elements comes in when activation must spread from the chunk node to the elements in order that the elements can themselves be expanded. So there should be an effect of number of elements in a cognitive unit only if it is necessary to inspect the content of these elements.

Establishing Correspondence with Old Knowledge

The key process in unitization is establishing a correspondence between a unit already in memory and some set of elements in working memory. The establishment of such a correspondence is far from a trivial process. The simplest case would be when the elements in working memory correspond exactly to the elements in the established knowledge unit. For instance, when we hear *George Washington crossed the Delaware*, presumably this proposition (and perhaps the exact sentence) corresponds to what we already have stored in memory. In this case it would not even be necessary to create a new unit. We could tag the existing memory unit. The many "chunking" proposals (Miller, 1956; Simon, 1974) advanced this

role for old knowledge in encoding new knowledge. However, it is often the case that the elements in memory do not correspond exactly to the old knowledge unit.

Some potential for the lack of correspondence between old and new was foreseen in the use of variables in the various schemata proposals. Consider the following story:

Fred went to the movies.
Fred had a headache.
He bought popcorn and a drink.
Fred saw *Star Wars*.
After the movie he went to the coffee shop.

Presumably, this does not correspond exactly to anything that we have heard before, but it can be put into correspondence with a schema of the following sort:

= Person goes to movies
= Person enters to = theatre
= Person pays for = ticket
= Person buys = refreshments
= Person gives = ticket to = collector
= Person finds = seat
= Person watches = movie
After = movie, = person leaves = theatre
= Person goes to = food establishment

Here = person, = theatre, = ticket, = collector, = refreshments, = seat, = movie, = food establishment, are all variables which can match a range of objects. The entire schema summarizes a set of events, none of which are identical. The variables denote places where one event might differ from another. Actually, in line with earlier remarks, this schema is probably too large to be a single cognitive unit and is better thought of as a hierarchy of units. Figure 6 illustrates a possible hierarchy.

Clearly, the problems of making correspondences with this structure are not so trivial. First, there is the problem of deciding whether the constants in the to-be-encoded structure can be put consistently in correspondence with the variables in the schema (i.e., Fred in correspondence with = person, popcorn and a drink in correspondence with = refreshments, *Star Wars* in correspondence with = movie, and coffee shop in correspondence with = food establishment). Moreover, we see that only some of the elements in the knowledge-unit are instantiated in working mem-

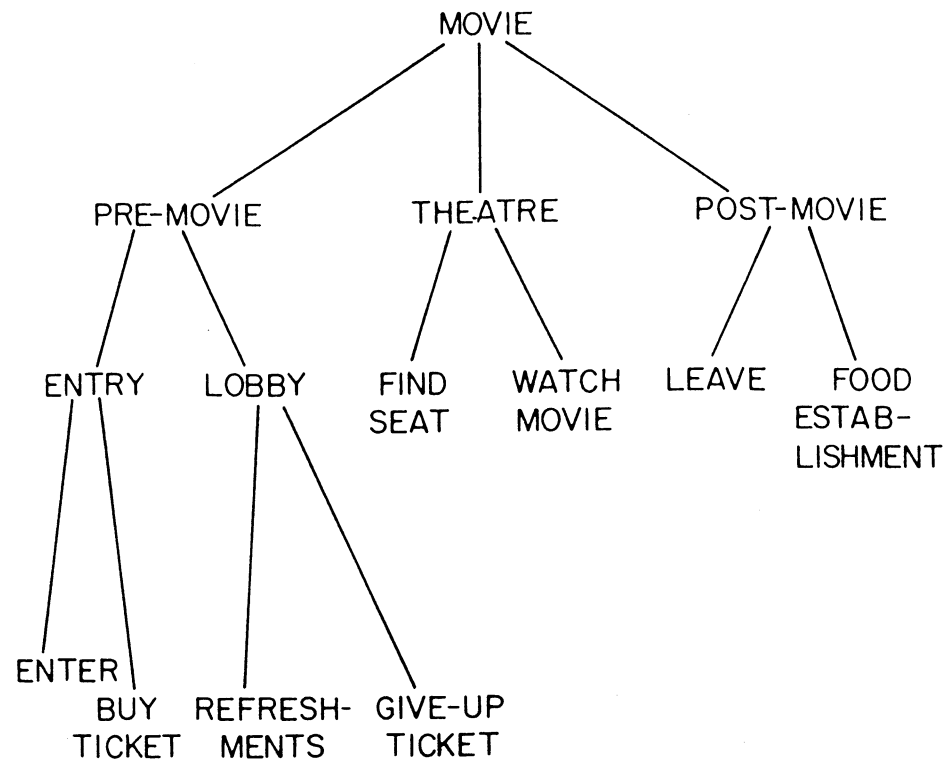


FIGURE 6. A representation of a possible hierarchy of units underlying the movie script.

ory. Moreover, there may be elements in working memory that do not correspond to the script (e.g., Fred had a headache). In its general form it is an extremely difficult matching problem to decide how to make the correspondence between such a variabilized pattern and the elements of working memory. A very good dissertation in computer science (Forgy, 1979) only solves part of the problem.

Related to the correspondence problem is the invocation problem. When should such a script be invoked. Presumably, when we hear "John bought popcorn and a drink," "John found a seat," we do not want to immediately invoke the movie script. Schank and Abelson make heavy use of certain key phrases like "went to the movies" to select a script. This seems related to the index in my proposal. However, this is only a partial solution.

For a long time we have suspected that the correspondence process might even be more complex than envisioned in the schema proposal with its variables and optional clauses. It seems that often correspondences are made with special cases rather than with general schemata. That is to say, oftentimes we spontaneously use

analogies to make correspondence, Consider the following:

Ahero in Imaginaryland

This story is about Ahero who led his nation of Imaginaryland through a bitter civil war because of his belief in basic equality of all sexes. Women in the eastern states of Imaginaryland were bought and sold as wives and had few basic rights guaranteed. The economy of the East was heavily agricultural and some farmers had acquired hundreds of wives to work on their huge farms. The practice had been abolished in the West because of the decline of agriculture and greater contact with the liberal ideas of Europe. Ahero grew up in the southern wilderness area of Imaginaryland. He developed a strong moral commitment to abolish the wife-practices in eastern Imaginaryland. After a hard political struggle Ahero became prime minister of Imaginaryland. He immediately abolished wife-trade and guaranteed all women rights almost equal to those of men. After a few months of abortive parliamentary moves, the eastern states seceded from the confederation. There followed a long bitter civil war which finally reunited Imaginaryland. Ahero was assassinated just before the end of the war by a western actor sympathetic to the eastern cause.

Presumably, a reader of this story finds herself inevitably led to making correspondences with Abraham Lincoln and the Civil War, although the correspondences are hardly perfect. However, finding correspondences here differs from the schema-correspondence in at least two ways: (1) The knowledge structure being evoked, for instance about Abraham Lincoln, is not schematic. In no way could it be considered to contain variables which are filled by constants in making the correspondence. (2) There are failures of correspondence—for instance, Ahero abolished wife-practice before the war, while Lincoln abolished slavery at the end of the war. Still the correspondences are made almost irresistibly. It seems that such correspondence-by-analogy can be frequently involved in understanding texts, and we suspected that it might have beneficial effects on memory like those associated with other manipulations that enable subjects to unify information by making correspondence with past knowledge.

In a recent paper (Schustack & Anderson, 1979), we have explored the issue of the effects of such analogies on memory. We used texts such as the following which do not have such obvious analogies:

Yoshida Ichiro was a Japanese politician of the 20th century. He was chosen as national leader in his own right after having substituted for another. He was responsible for intensifying his country's involvement in a foreign conflict. He devoted much money and effort to eradicating economic and social injustices. He was made aware of rapidly worsening problems with his nation's energy supply, but he ignored them out of fear of political repercussions.

We contrasted four conditions—an experimental condition and three kinds of control. In the experimental condition the subject was told of the analogy between the fictional person and a well-known individual (in the case above, Lyndon Johnson). Such analogies were deliberately not perfect but are fairly close. Using a norming study, we had designed the material so that subjects would not spontaneously discover the analogies. In the first control condition subjects were not given any information about an analogy. The other two control conditions were used to assess the effects of having a name given as an analogy when no correspondence could be made. In the counterfeit condition subjects were given as an analogy a name of a person they did not know (e.g., Michael H. Donovan). The final condition, the inappropriate condition, was to test for the effect of providing as an analogy the name of a famous person. Here we provided an analogy that did not fit the fictional person well (e.g., Robert F. Kennedy for the Yoshida Ichiro character). The result was that subjects showed by far and away best memory for the material when they were given an appropriate analogy. This seems to support our suspicion that subjects find themselves able to establish loose analogical correspondences with past knowledge and that these correspondences really can help memory. The analogy helps because it allows the subject to interlink all the elements for which correspondences have been made in a single step rather than having to separately encode each element.

Thematic Relatedness

Central to the discussion has been the idea of a correspondence between incoming data and existing knowledge structures. I have been arguing that such correspondences will be made and will have their impact on memory even in cases where the correspondence is hardly perfect. There are various kinds of "imperfection" that we have considered so far: (1) Certain pieces of the incoming data may

not be matched in the existing structure and must be "skipped over." (2) Certain elements of the incoming data may be mismatched in the existing structure and require either the use of variables or of analogy. These are examples of where the data do not quite match the existing knowledge structures. There is another possibility, which is that no existing knowledge unit may adequately match a set of data elements, but that the combination of a number of existing knowledge units will match the data. This is what happens in the case of thematic relatedness. As we noted with respect to the stories like the one represented in Figure 1, these appeared to be unitized, although they corresponded to no single preexisting knowledge structures. Similarly, Reder and Anderson (in press) found the same effects for sets of facts that were instantiations of familiar sequences (i.e., scripts) versus sets which were only loosely and thematically related.

It remains unclear what exactly it takes for a set of facts to become thematically related. At one extreme the facts might only need to be related by interassociated structure. At the other extreme, the facts may need to be much more tightly interlocked by some plan. Schank and Abelson (1977) have considered such possibilities in their discussion of plans and themes. We have yet to do the research that will determine just how loose the thematic structure can be before it loses the ability to integrate incoming information.

I see the problem of computing this correspondence between existing knowledge and incoming data as part of a major issue for cognitive science: that of computing partial matches. A presupposition underlying information-processing work from Sternberg's analysis of memory scanning to production systems is that we detect perfect matches between structures or mismatches. However, it is perfectly clear that the human system operates in terms of partial matches, and any viable artificial system will have to also. The closest thing we have to this in psychology is the work on random-walk models of recognition (Link, 1975; Ratcliff, 1978); however, these are not up to the structural complexity of the input that need to be matched. We are currently working on a "partial" solution to this partial-match problem. As these ideas are not yet completely implemented, I don't think it appropriate to present them yet. However, until they are presented, there is a very big hole in the theory presented here. The basic claim of this paper is that the act of correspondence creates a cognitive unit. To have a real theory, then, one needs to specify the partial-match processes by which this correspondence is computed.

COMPARISON WITH OTHER THEORIES AND CONCEPTS

It is useful to consider how this proposal relates to other proposals that have been offered for knowledge representation and memory. The past proposals for which it is easiest to establish a correspondence are the ideas about how higher order concepts could serve to organize instances in a list. This is the work on categorized recall (e.g., Bower, Clark, Lesgold, & Winzenz, 1969; Mandler, 1967; Tulving, 1968). In this work one finds the proposals for a hierarchy of units, for all-or-none retrieval of units, and for the hypothesis that heightened recall results from unitization. These are all ideas that I want to emphasize in my proposal.

There are two principal differences between this proposal and standard discussions of categorized recall. One, of course, is the extension to propositions and other larger units of knowledge. The second is the explicitness of the mechanisms proposed. It has always been something of a mystery just why categorization produces the effects it does. If it produces better recall because it takes advantage of existing knowledge structures, as was commonly supposed, it becomes a mystery why there are not more intrusions of nonpresented members of the category. One solution to this mystery was the tagging proposal developed in the context of the FRAN model (Anderson, 1972; Anderson & Bower, 1972). This tagging proposal left open the issue of why it was easier to associate tags to elements in existing structures rather than simply interassociating them in new structures. While not denying the tagging mechanism, a rather different mechanism is being proposed here. The proposal here is that a new unit is created that organizes whatever elements are in working memory. The problem for this mechanism is to explain why there are any categorical intrusions rather than why there are so few. Perhaps this could be explained in terms of a guessing strategy.

The ideas in this paper have their strongest superficial similarity to the set of ideas that have developed in the past five years about ~~frames, scripts, and schemata, and I have made constant reference~~ to these ideas throughout the paper. There is a range of ideas that have been advanced in these other proposals, and I would not want to advance any aspect of my proposal as totally unique; but I would like to emphasize certain aspects that are relatively distinct: (1) Important is the concept of a hierarchy of units and the idea that the need for this hierarchy is forced by limits on the size of working

memory. (2) Related to (1) is the conception of a cognitive unit as appropriate at many levels including the proposition and large story texts. (3) Also important is the extremely general notion of *correspondence* being advanced. Correspondences are not just calculated by fitting constants into variable slots but rather by a more general process of partial matching. (4) A fourth relatively unique aspect of the proposal is the explicitness with which the consequences of the cognitive-unit idea have been developed for mechanisms of encoding and retrieval (i.e., we have described the mechanisms that lead to all-or-none recall and heightened levels of recall.)

Another recent proposal, to which mine bears a superficial similarity, is the "cogit" model of Hayes-Roth (1977). In that model she proposed that elements could become unitized and that once unitized they would acquire an immunity to interference. The similarities here are quite superficial, and the two proposals are really quite orthogonal. It is important to state why they are orthogonal. First, Hayes-Roth proposed that the elements became unitized by dint of massive practice, while the proposal here is that unitization occurs after a single study when the elements are put in correspondence with an existing knowledge structure. Second, Hayes-Roth claimed that once knowledge was unitized in her sense, the knowledge unit became impervious to interference. When elements are unitized in the sense of this paper, the consequence for interference in many tasks is that the individual elements within the unit are no longer interfering or interfered with. However, the knowledge unit in which they are unitized still interferes and can be interfered with. These remarks are not meant to question the usefulness of Hayes-Roth's ideas, but only to question perceptions of their relevance to the current discussion.

The current proposal might seem to be in contradiction with another proposal I have put forth to account for the beneficial effects of prior knowledge in the acquisition of new knowledge. This is the elaboration proposal (Anderson, 1976; Anderson & Reder, 1979; Reder, 1979). The idea is that subjects generate additional redundant memory connections to encode a particular proposition. For instance, a subject given the sentence *The doctor hated the lawyer* might elaborate it: *The lawyer brought a lawsuit against the doctor for malpractice. When the doctor was on the witness stand, the lawyer assailed him with a stream of questions. The doctor glared at the lawyer. The lawyer continued to ask insulting questions. The doctor cursed the lawyer.* A subject, even if he could not remember the original

sentence, might be able to infer it if he could retrieve a substantial amount of his elaboration. The elaboration proposal is that subjects use past knowledge to create inferential redundancies that help them remember new information. The unitization proposal claims that past knowledge helps retention of new information because the new information can be more rapidly encoded in terms of its correspondence to past knowledge. Thus one proposal emphasizes redundancy of encoding, while the other emphasizes increased rate of encoding.

While the two proposals sometimes can both account for the same memorial advantage of familiar over unfamiliar material, they extend to explain different phenomena. It is hard to make any firm connections between the elaboration proposal and any unitization phenomena such as all-or-none recall or within-unit priming. On the other hand, the unitization proposal cannot account for the effect of embellishment when these embellishments do not serve to establish correspondences with prior-knowledge units. For instance, it is hard to see why the doctor-lawyer elaboration given earlier would help according to the unitization proposals. Related to this is the fact that it is hard to extend the unitization proposal to account for effects of depth of processing or inferential intrusions in recall (see Anderson & Reder, 1979). So, my inclination is to consider both hypotheses as necessary to account for the complexities of human memory.

SUMMARY

In this paper I have reviewed the data and theoretical considerations that bear on the issue of whether propositions and schemata should be considered cognitive units. Assuming a certain general framework about working memory and long-term memory, the evidence tended to point toward a positive conclusion—that both can be cognitive units. Evidence for this comes from observations about all-or-none recall, heightened recall of units, associative priming, diminution of interference effects, considerations of implementation efficiency, and considerations of cognitive naturalness. This being said, I do not find the current picture so convincing that I would want to make a never-say-die commitment to the positive conclusion for cognitive units at all levels of knowledge structures. The idea seems sufficiently promising that it is worth-

while to develop a more explicit theory of what a cognitive unit would be like. So I have proposed a general notion of a cognitive unit that spanned propositions and schemata as special cases and specified the encoding and retrieval properties that such a cognitive unit would have. To be succinct, the important ideas associated with cognitive units are the following: (1) They can occur at multiple levels and enter into hierarchies. (2) If a set of working-memory elements can be put in correspondence with an existing knowledge structure, the elements can be joined in a cognitive unit by a single encoding act. (3) The elements in a cognitive unit are brought into working memory in a single retrieval act. (4) It is possible to evaluate general properties of a cognitive unit without having to expand it into its elements and inspect these. Perhaps this is done by measuring level of activation.

REFERENCES

- Anderson, J. R. FRAN: A simulation model of free recall. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 5). New York: Academic Press, 1972.
- Anderson, J. R. Retrieval of propositional information from long term memory. *Cognitive Psychology*, 1974, 5, 451-474.
- Anderson, J. R. *Language, memory, and thought*. Hillsdale, N.J.; Lawrence Erlbaum Associates, 1976.
- Anderson, J. R. Arguments concerning representations for mental imagery. *Psychological Review*, 1978, 85, 249-277.
- Anderson, J. R. Further arguments concerning representations for mental imagery: A response to Hayes-Roth and Pylyshyn. *Psychological Review*, 1979, 86, 395-406.
- Anderson, J. R., & Bower, G. H. Recognition and retrieval processes in free recall. *Psychological Review*, 1972, 79, 97-123.
- Anderson, J. R. & Bower, G. H. *Human associative memory*. Washington, D.C.: Hemisphere Press, 1973.
- Anderson, J. R., & Paulson, R. Representation and retention of verbatim information. *Journal of Verbal Learning and Verbal Behavior*, 1977, 16, 439-451.
- Anderson, J. R., & Reder, L. M. An elaborative processing explanation of depth of processing. In L. S. Cermak & F. I. M. Craik, (Eds.), *Levels of processing in human memory*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1979.
- Anderson, R. C. Substance recall of sentences. *Quarterly Journal of Experimental Psychology*, 1974, 26, 530-541.
- Black, J. B., & Bower, G. H. Episodes as chunks in narrative memory.

- Journal of Verbal Learning and Verbal Behavior*, 1979, 18, 309-318.
- Bobrow, D. G., & Winograd, T. An overview of KRL, a knowledge representation language. *Cognitive Science*, 1977, 1, 3-46.
- Bower, G. H., Black, J. B. & Turner, T. J. Scripts in memory for text. *Cognitive Psychology*, 1979, 11, 177-220.
- Bower, G. H., Clark, M. C., Lesgold, A. M., & Winzenz, D. Hierarchical retrieval schemes in recall of categorical word lists. *Journal of Verbal Learning and Verbal Behavior*, 1969, 8, 323-343.
- Broadbent, D. E. The magical number seven after fifteen years. In R. A. Kennedy & A. Wilkes (Eds.), *Studies in long-term memory*. New York: Wiley, 1975.
- Cohen, B. H. Recall of categorized word lists. *Journal of Experimental Psychology*, 1963, 66, 227-234.
- Cohen, B. H. Some-or-none characteristics of coding behavior. *Journal of Verbal Learning and Verbal Behavior*, 1966, 5, 182-187.
- Forgy, C. L. *On the efficient implementation of production systems*. Ph.D. Dissertation, Computer Science Department, Carnegie-Mellon University, 1979.
- Goetz, E. T., Anderson, R. C., & Schallert, D. L. *The representation of sentences in memory*. Unpublished manuscript, 1979.
- Hayes-Roth, B. Evolution of cognitive structures and processes. *Psychological Review*, 1977, 84, 260-278.
- Hendrix, G. G. *Expanding the utility of semantic networks through partitioning*. Stanford Research Institute Technical Note 105, 1975.
- Katz, J. J., & Fodor, J. A. The structure of a semantic theory. *Language*, 1963, 39, 170-210.
- Kintsch, W. *The representation of meaning in memory*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1974.
- Landauer, T. K. Memory without organization: Properties of a model with random storage and undirected retrieval. *Cognitive Psychology*, 1975, 7, 495-531.
- Link, S. W. The relative judgment theory of two choice response times. *Journal of Mathematical Psychology*, 1975, 12, 114-135.
- Mandler, G. Organization and memory. In K. W. Spence & J. A. Spence (Eds.), *The psychology of learning and motivation* (Vol. 1). New York: Academic Press, 1967.
- Mandler, J. M., & Johnson, N. S. Remembrance of things parsed: Story structure and recall. *Cognitive Psychology*, 1977, 9, 111-151.
- Miller, G. A. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 1956, 63, 81-97.
- Minsky, M. A framework for representing knowledge. In P. H. Winston (Ed.), *The psychology of computer vision*. New York: McGraw-Hill, 1975.
- Newell, A. Production systems: Models of control structures. In W. G. Chase (Ed.), *Visual information processing*. New York: Academic Press, 1973.

- Norman, D. A., & Rumelhart, D. E. *Explorations in cognition*. San Francisco: W. H. Freeman, 1975.
- Owens, J., Bower, G. H., & Black, J. B. The "soap opera" effect in story recall. *Memory and cognition*, 1979, 7, 185-191.
- Pylyshyn, A. W. Validating computational models: A critique of Anderson's indeterminacy of representation claim. *Psychological Review*, 1979, 86, 383-394.
- Ratcliff, R. A theory of memory retrieval. *Psychological Review*, 1978, 85, 59-108.
- Ratcliff, R., & McKoon, G. Priming in item recognition: Evidence for the propositional structure of sentences. *Journal of Verbal Learning and Verbal Behavior*, 1978, 17, 403-417.
- Reder, L. M. The role of elaborations in memory for prose. *Cognitive Psychology*, 1979, 11, 221-234.
- Reder, L. M., & Anderson, J. R. A partial resolution of the paradox of interference: The role of integrating knowledge. *Cognitive Psychology*, 1980, 12, 447-472.
- Rumelhart, D. E. Understanding and summarizing brief stories. In D. La Berge & J. Samuels (Eds.), *Basic processes in reading: Perception and comprehension*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.
- Rumelhart, D. E., & Ortony, A. The representation of knowledge in memory. In R. C. Anderson, R. J. Spiro, & W. E. Montague (Eds.), *Schooling and the acquisition of knowledge*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1976.
- Rychener, M. D., & Newell, A. An intractible production system: Basic design issues. In D. A. Waterman & F. Hayes-Roth (Eds.), *Pattern-directed inference systems*. New York: Academic Press, 1977.
- Schank, R. C. Conceptual dependency: A theory of natural language understanding. *Cognitive Psychology*, 1972, 3, 552-631.
- Schank, R. C. The structure of episodes in memory. In D. G. Bobrow & A. M. Collins (Eds.), *Representation and understanding*. New York: Academic Press, 1975.
- Schank, R. C., & Abelson, R. P. *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.
- Schustack, M., & Anderson, J. R. Effects of analogy to prior knowledge on memory for new information. *Journal of Verbal Learning and Verbal Behavior*, 1979, 18, 565-583.
- Simmons, R. F. Some semantic structures for representing English meanings. In R. Freedle & J. B. Carroll (Eds.), *Language comprehension and the acquisition of knowledge*. Washington, D.C.: Winston, 1972.
- Simon, H. A. How big is a chunk? *Science*, 1974, 183, 482-488.
- Smith, E. E., Adams, N., & Schorr, D. Fact retrieval and the paradox of interference. *Cognitive Psychology*, 1978, 10, 438-464.
- Thorndyke, P. W. Cognitive structures in comprehension and memory of narrative discourse. *Cognitive Psychology*, 1977, 9, 77-110.

- Thorndyke, P. W., & Bower, G. H. Storage and retrieval processes in sentence memory. *Cognitive Psychology*, 1974, 5, 515-543.
- Townsend, J. T. Issues and models concerning the processing of a finite number of inputs. In B. H. Kantowitz (Ed.), *Human information processing: Tutorials in performance and cognition*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1974.
- Tulving, E. Theoretical issues in free recall. In T. R. Dixon & D. L. Horton (Eds.), *Verbal behavior and general behavior theory*. Englewood Cliffs, N.J.: Prentice-Hall, 1968.